

N91-20721

ISSUES ON COMBINING HUMAN AND NON-HUMAN INTELLIGENCE

Irving C. Statler
NASA Ames Research Center
N239-1
Moffett Field, CA 94035

Mary M. Connors
NASA Ames Research Center
N262-5
Moffett Field, CA 94035

Introduction

For the foreseeable future, there will be very few activities or missions that will be accomplished entirely by non-human, totally autonomous systems. Human intelligence and the ability it confers to exercise judgment and, thus, deal with unexpected situations will warrant the services of human members in future systems. The number of autonomous systems working in conjunction with, or in support of, human crews has been growing rapidly and can be expected to grow at an even faster rate in the future. We are faced with the problem of designing systems in which a machine intelligence and a human intelligence can work together as partners. This may be more difficult than designing a fully automatic, unmanned system. Unfortunately, we have little appreciation of either the potential or the limitations of close working relationships between humans and intelligent machines, or of how these interactions affect relations with other crew members or total crew performance.

The purpose of this paper is to call attention to some of the issues confronting the designer of a system that combines human and non-human intelligence. We do not know how to design a non-human intelligence in such a way that it will fit naturally into a human organization. Our concern is that, without adequate understanding and consideration of the behavioral and psychological limitations and requirements of the human member(s) of the

system, the introduction of artificial intelligence (AI) subsystems can exacerbate operational problems. We have seen that, when these technologies are not properly applied, an overall degradation of performance at the system level can occur. Only by understanding how human and automated systems work together can we be sure that the problems introduced by automation are not more serious than the problems solved.

Background

Our experience with automation in space is still quite limited. However, there are examples from aircraft operations to illustrate the point that the implementation of engineering "solutions" may prove inadequate when human behavior is involved. A number of incidents (Connors 1989, Wiener & Nagel 1988) have raised questions about our ability to combine humans and automation into effective teams. Although we will be referring here primarily to examples from aircraft cockpits, the problems of man-machine integration in complex systems are ubiquitous. It is easy to strike out "cockpit" and fill in "air traffic control center", "submarine", "nuclear power plant", "launch control center", "space station", or "Mars vehicle". For many years, we have been able to rely on the adaptability of the human to take maximum advantage of each new technology. In the current environments of data-display that missions have required and computers have enabled, our man-machine systems are capable of saturating the human component with the sheer number of

displays to be read, controls to be engaged, and decisions to be made. Nevertheless, we continue to depend on the human pilot to assess the situation instantaneously, to make the "right" decision, and to initiate the appropriate action. Many of us believe that, in the realm of both military aircraft and space systems, we are close to the practical limitations of human sensory and cognitive capabilities.

For example, over the years, electro-mechanical instruments, switches, and buttons have propagated wildly in the cockpit, filling all the available space. When the designer of the modern fighter aircraft cockpit was faced with the dilemma of reduced display space in the smaller cockpits along with the need for still more information to be displayed to the pilot, his solution was to replace task-specific displays and controls with multi-purpose displays and multi-function controls. Although this solution addresses the narrowly-defined display problem, it does not solve the operational problem, since modern computers that are brought aboard to drive these displays are capable of presenting far more data than a human can possibly access and assimilate in real time.

The F-18 aircraft has one of the more advanced cockpits and is a good example of the problem of data overload. This cockpit has three cathode-ray tubes and a head-up display. There are 675 acronyms and 177 symbols that can appear in four different sizes on any of the three cathode ray tubes. There are 73 threat, warning, and caution indicators, 59 indicator lights, and 6 warning tones (no messages, just tones), 10 multi-function switches on the throttle, 7 on the stick, 19 controls on the panel underneath the head-up display, and 20 controls around the periphery of each of the three cathode-ray tubes, each of which has a multi-switch capability. Most of the data displayed requires that the pilot's foveal vision be engaged (while peripheral vision, utilized in earlier displays, is largely ignored.) Every piece of information that is available to the pilot for multi-purpose display requires an additional control to access that information. This imposes a memory load on the crew who must remember how to access the desired information and how to perform the

required control function. Often, these controls must be found and actuated by touch while the pilot is visually engaged elsewhere, sometimes during moments of extreme physical and mental stress. Since not all of the information about his aircraft can be displayed to the pilot at all times, there has evolved a proliferation of warning and alerting systems. These systems remind pilots to take actions, call attention to deviations from expected ranges, suggest or demand an action, warn of unacceptable configurations, and even take action on their own.

One of our favorite examples of where the engineering solution to a problem seems to disregard basic human-factors principles is the helmet-mounted display for the US Army's attack helicopter called the Apache. When a military helicopter is operating nap of the earth at night or in adverse weather, the pilot desperately needs help. He must be able to see something of the outside world. There is an infrared sensor, called a FLIR, in the nose of the helicopter that provides a display in the cockpit. There is also a computer on board that generates symbologies both for flight-control information and for weapons-control information. There are 19 such symbols in three different formats, depending on the flight phase. For the pilot of this aircraft, all of this (the FLIR image with superimposed flight-control symbologies and weapon-control symbologies) is presented on a two and a half centimeter monocular over his right eye. At the same time, his left eye is expected to take care of the contextual scene and the instrument panel.

We must also keep in mind that equipment intended to enhance human capability can actually encumber it by exacting a physiological toll that, in turn, compromises performance. The tendency to attach devices to the head is of particular concern, often leading to a loss of head mobility and fatigue. We have found some equipment of this type to cause both physiological and psychological problems in our military pilots. For example, some current helmet-mounted displays provide different and potentially disorienting visual images to the two eyes.

The typical military or civil pilot today must integrate enormous amounts of data from many dissimilar sources, sometimes under great time pressure. In our attempts to maximize the number of physical channels available for transferring these data, we have introduced voice and other aural displays. However, the addition of secondary modalities does not double the human's information processing capability; indeed it may even impede it by distracting the operator at a critical time. In fact, the operator may not even be aware of additional information because humans, under certain conditions, tend to narrow their attention. The problem may be further exacerbated by the human tendency in stressful situations to see what he expects to see and to hear what he expects to hear. Both the civil and the military sectors provide examples of where warning signals have been ignored due to the human tendency, when under stress, to selective attention.

The main point we wish to make is that humans, although highly adaptable, are not unlimited in their ability to accommodate to demanding task environments. In some of our more complex cockpits, the human may no longer be able to "take up the slack". In addition, the electronic systems we are now providing to aid the pilot may not be helping at all, and may actually be complicating his job. He is confronted with too much data and in formats that may not be conducive to rapid interpretation. It is useless to continue providing more data if the operator is unable to use it, since it is relevant information, not data, that is needed if the operator is to make good decisions.

In the past, when similar situations have been encountered, we have typically solved the problem by putting more men on the job. There are many situations in which this solution is impractical, and so it is tempting to look to artificial intelligence (AI) as a way of augmenting human capabilities. Presumably, with AI, one could fuse sensor outputs, integrate data, present only what was needed when it was needed, and assist the pilot in making decisions.

In keeping with this view, there have been proposals for military aircraft with one human pilot and several electronic crew members. Artificial intelligence, decision-support systems, knowledge-based systems, and expert systems became the buzz words of the eighties. However, although there has been a great deal of casual talk about the role that machine intelligence might play, the problem of developing the essential symbiotic relation between human and non-human intelligence has been examined only cursorily. We really do not understand what it takes to satisfy human needs, and it appears that even if we did, we do not yet know how to build it.

The Problem

Knowledge-based and expert systems have found some limited application in the control of physical plants, manufacturing processes, and quality control. However, they have yet to find a role in circumstances that cannot be described with mathematical algorithms or logical rules. But, not all knowledge is susceptible to logic.

There exist many potential applications for knowledge-based systems. Unfortunately, there are several fundamental things that we still do not know how to do. Following are just a few:

1. how to develop the complete knowledge base (or even know when or if it is complete,) particularly if it does not lend itself to logical rules;
2. how to have an expert system learn from experience by changing its rules;
3. how to enable the system to make complex decisions in real time during unexpected situations;
4. how to assure compatibility with the human operator's perceptions of the situation and acceptability by the operator of recommended solutions; and
5. how to validate the "sanity" of the system.

As AI grows and progresses, we can expect some advances in knowledge and understanding of these areas. However, automated systems will remain

limited by the assumptions that created them, i.e., they will always be "blind" to conditions that were not explicitly or implicitly included in their design (Winograd and Flores, 1987).

Also, while computers can, after a fashion, think and learn, they do not think or learn as humans do. Consequently, if computational systems should take over decision-making chores, the human operator may find himself at odds either with what the computer is doing or the way in which it is doing it.

The rationale behind the introduction of automation has been the desire to enhance total system capabilities while maintaining operator workload at acceptable levels, thereby minimizing the possibility of human error. However, as more and more physical control activities have been successfully automated, they have been replaced by mental activities on the part of the human operator. Our experience with automation indicates that its introduction usually relocates and changes the nature and consequences of human error, rather than removing it.

The negative reactions and incident reports that NASA is beginning to receive regarding the electronic crew member in the glass cockpit of our modern civil transports support our concern. The glass cockpit has been criticized for its failure to reduce mental workload. Pilots believe that automatic devices demand constant attention and each device creates its own demands on the pilot's time. Automation tends to isolate the flight crew from the state of the aircraft and the modern pilot can feel not only left out of the loop, but externally controlled. Recent accidents suggest that excessive automation tends to lower the level of vigilance of human operators. Moreover, automation frequently addresses short-term, subsystem solutions, rather than total system performance. For instance, there is often inadequate feedback and interaction with the human controller (Norman 1990). Consider, for example, the system that corrects for a fault without notifying its human partner. In one incident, a race car equipped with the latest automatic compensation for brake failures suffered a failure

in one brake. The system automatically compensated, just as it was designed to do. Shortly after, a second brake failed, and, due to the increased loading, the third-brake failure quickly followed the second. But, the automatic compensation system had done its job so well that it was not until the fourth brake failed that the driver realized he had a problem. This is an example of a faulty design philosophy that has as its goal to show the operator only what he needs to know when (someone else determines) he needs to know it.

In an analogy with the artificial heart program, the introduction of AI in a given system can fail (and has failed) because we do not understand the reaction mechanisms of the human. An AI subsystem must be designed to sing and dance gracefully with the human crew as well as with the energy sources that power it and the environment in which it must operate.

Therefore, the total system design must take into account the capabilities, limitations, and needs of the human component. We do this already with respect to human physiological constraints, but now we must take into account cognitive, motivational, and other psychological needs. We will continue to rely on the human in the vehicle for creativity and innovation in coping with the unexpected. In our future space systems, these humans will be better trained and more knowledgeable than ever before; but they remain humans whose tolerance for vibration, heat, hypoxia, and G-forces has not changed; whose visual perception and information-processing capacity are still limited; and whose decision-making ability remains susceptible to fatigue, illusions, biases and stress.

The design of the equipment intended to improve total system performance must consider the full impact it has on human behavior and on the human's ability to perform the role expected of him. This requires consideration of such things as the effects on humans of being "in the loop" or "out of the loop", the nature of trust between humans and machines, the ability of the machine to communicate the reasons for its actions to the

satisfaction of the human operator, the ability of the machine to respond to the human's "what if I did it this way?" queries (Galdes and Smith, 1990) and the fact that the human needs to feel that he or she is ultimately in control. How can we be certain that any data display will be clear and unambiguous in all situations, so as to ensure the correct interpretation by the human for fast and accurate reaction in the rare critical situation? How do we keep the human well informed without annoying him? If the machine carries out all the routine tasks, how is the human to be kept in a state of alertness in which he or she is capable of performing adequately if the machine should fail? Many decisions regarding whether or not to manually override an automatic system will need to be made during critical phases of missions. Given the demands of these phases, does the automated system provide a net benefit to the crew? Can the workload required of the human crew during these periods be kept within acceptable limits?

Involving the human in the decision-making process provides an essential layer of checks and balances to make up for the shortcomings of the non-human intelligence. However, there is no point in extolling and relying on the real or imagined virtues of human creativity and innovation if the human doesn't know when to take control, or if the system design is such that the human is unable to be creative or innovative in the actions which the system allows him to initiate.

Often the problem of the human-machine interaction is considered to be merely one of interface design. This viewpoint is a dangerous oversimplification. It is like suggesting that human communication can be explained on the basis of word recognition. System functionality depends on characteristics of the communicating systems that extend well beyond issues of the operator interface. AI is going to be used to support dynamic interactive tasks in which the human mind is an important and active component of the total system. Designing tools for this kind of complex cognitive-psychological activity goes well beyond the issue of

display and control interfaces. It can no longer be viewed as a process of designing a machine to do something, and then designing the information displays and controls which enable the operator to guide the machine. A system's usability is determined by the details of a given design and not just by its interface style.

Approach: The Crew System

The introduction of the concept of artificial intelligence to work with the human requires that we begin to think, not in terms of a human operating a machine as we have in the past, but in terms of communication between intelligent agents. The problem of designing a system that produces a symbiotic integration of the powers of the human brain and computers is incredibly complex and difficult. It is not simply a question of the proper allocation of functions between man and machine, nor should the human and the machines be considered in competition for duties. Rather it is essential that the human and the machine are explicitly considered as parts of a larger functioning system. The human may no longer be the sole supplier, as in the past, of the initiative, the direction, the integration, and the standards. For instance, it may be that the safest and most efficient system will be one that incorporates considerable duplication or interchangeability of functions among its human and non-human crew members and thus benefits from the strengths of both. A joint cognitive system implies a productive relationship between the knowledge of the machine and that of the human in which the different points of view are integrated in the decision process.

In a previous paper, one of us used the term "crew system" to describe all active, intelligent flight participants, whether human or artificial. Dr. Malin at JSC has proposed the idea of making humans and computers "team players". The implication of these terms is that the human(s) and the machine(s) must be considered as forming a partnership, sharing all the responsibilities and authorities in a concept of cooperation rather than one of human or machine control. The close

coupling of humans and machines requires us to view their interactions as a total system design problem; i.e., a crew that is composed of both human and non-human intelligence.

One requirement of this integrated-design concept is for training and support to help humans cope with the new electronic environment. A second, and more pressing requirement, is to learn to design machine components for compatibility with real human behavior and with full recognition that human beings experience fluctuating motivation and attention and also make errors.

System design geared to blending human and automated systems must take into account all levels of human activity from the most basic perceptual response, through man-machine interface, and up to and including full integration into the relevant environment. For a human to perform a particular task, he must be able to translate his psychological representations of the system state, his goals, and his intentions into physical actions. To interpret the outcome of his actions, the human must be able to perceive the resulting system state and relate those perceptions to his psychological representations. We must understand how people recognize patterns, integrate information, add their own previous knowledge and value structure and come up with intelligent, appropriate decisions under difficult circumstances. A problem will ensue if the non-human intelligence negatively interferes with any part of this fundamental process (Norman 1987).

The need for considering design from the aspect of a crew system also introduces concerns related to small group and organizational science. We need to expand our view of system requirements to include information processing and motivation of multiple agents in organizations. When we introduce a non-human intelligence into the crew, the entire interactional structure of the crew changes. At these higher levels of integration, the results of NASA's extensive research in group dynamics pertaining to flight crews of long-haul civil air transports are particularly relevant. For example,

in human groups it has been found that junior members are often reluctant to question the actions of the senior member even in critical situations. Similarly, automated systems that are perceived as highly reliable or having a high level of authority have produced an unwillingness on the part of the human to question and override. The quality of interpersonal interactions and coordination among the members of a crew in terms of their behavior and communications has been shown to be a fundamental factor in the performance of that crew and its susceptibility to errors. For human crews, this problem is a matter of selection, training, and organizational management; for the non-human member, it is a matter of design; for the entire system, all these factors, along with integrating procedures, must be included.

As yet, the human factors community has been unable to consolidate its empirical data into design methods and principles to guide the design process. The demands for performance-enhancing human/automation systems exceed our present understanding of the science. There are too many uncertainties in what principles are relevant to what tasks; empirical emphasis tends to be placed upon isolated properties of individual processes; and even well-established phenomena developed in laboratory settings often have very different levels of influence when imbedded in more complex tasks.

Since comprehensive design guidelines have been unavailable, system developers have attempted to assess the qualities of systems composed of AI and human components in a post facto manner. Thusfar, the index of acceptability has tended to be that the AI system has reached operational status. This is an unacceptable validation procedure and begs the question of total system capability. New indices of quality and acceptability are needed and even basis assumptions should to be re-examined.

In considering what might happen in combining human intelligence and artificial intelligence, one might postulate four major outcomes: (1) performance (in terms of effectiveness, efficiency, cost, etc.) is equal to that of the human crew alone;

(2) performance is equal to that of the automated system; (3) performance is less than that of the human crew or of the automated system alone, and (4) performance is better than either system alone. In general, only the fourth outcome (improved system performance) justifies the investment required for combined systems. The task then becomes finding practical methods and appropriate metrics for assessing the level of performance and the facility with which the human and the machine cooperate to solve unexpected problems. This task represents a substantial challenge to both the AI and the human factors communities.

A paper presently in preparation by one of the authors (Connors and Harrison, 1990) outlines research issues that are likely to be important in combining human and non-human intelligence. As this paper points out, one way to begin to understand the possibilities of integrated systems is to fully understand the failures that occur in present systems. A useful approach is to analyze the specifics of how human error changes (if at all) in the presence of automated systems. It is not enough, however, to examine error events in terms of number, severity, point in the mission, and the like. Critical information may be lost if one fails to examine error (or other measurable change) in terms of the human functions impacted (i.e., perception, recognition, attention, memory, information processing, coordination, and the like.) All opportunities, whether in simulation, field studies, or actual operations need to be utilized to begin to appreciate the dynamics of human behavior in human/automated settings. Cumulatively, this experience base could help focus future research and, eventually, to establish selection, training, procedural and design criteria.

Conclusion

Currently, systems are being planned based on exceedingly generous estimates of the human's capabilities for processing information and of the artificial intelligence capabilities for making sound decisions that are accepted by the human. In other words, we are busy building solutions when we do not yet fully understand the problem.

Without an understanding of how to combine human and non-human intelligence effectively, we shall be unable to implement rational designs for our future space systems. The issues we have raised here, and others, need to be examined when considering the potential of these systems. The need for an effective marriage of human and non-human intelligence will increase greatly with the advent of Space Station Freedom and with the subsequent, more distant, missions. Life in these space vehicles is likely to mimic life in other isolated and confined settings, i.e., marked by fatigue, moodiness, disturbed sleep, sensory deprivation, reduced motivation, and loneliness (Connors, Harrison and Akins, 1985; Harrison and Connors, 1984). All of these will tend to exacerbate the physical problems that the space crews will endure. Yet, the crew must not only survive, but display a high level of productivity. In the longer-durations space missions of the future, the use of automation and the discharge of responsibilities by human and non-human crewmembers will be essential to the conduct of the mission as well as to the health and welfare of the crew.

While we stress, as we do here, the problem of data and activity overload, we should keep in mind that, during some phases of long-duration spaceflight, the opposite problem may occur. Boredom during long and uneventful phases of flight could lead to loss of productivity and it may be necessary to design into these system a level of crew workload that is not only sufficiently restricted to be manageable, but also sufficiently large and engaging to offset boredom and ennui (Statler and Billings, 1989).

One day, some believe, the intelligence of a computer may rival that of the human brain. One day, we may learn how to couple human brains and computing machines in new and productive partnerships. For now, however, we must rely predominantly on human intelligence, judgement, flexibility, creativity and imagination in dealing with unexpected events; while relying heavily on machine intelligence for the logic, speed, persistence, consistency and exactitude it possesses.

Our task for the near future is to begin the process of building towards symbiosis and improved system performance, avoiding on the way, the pitfalls that could lead to precipitous system failure.

References

- Connors, Mary M. (1989) Crew System Dynamics: Combining Humans and Automation Presented at the 19th Intersociety Conference on Environmental Systems, 24-26 July 1989, San Diego. SAE Technical Paper Series 891530; 1989 Transactions in press.
- Connors, Mary M. and Harrison, Albert A. (1990) Combining Humans, Automation and Telecommunications in Space-Analog Environments. Paper in preparation.
- Connors, Mary M., Harrison, Albert A. and Akins, Faren F. (1985) Living Aloft: Human Requirements for Extended Spaceflight. NASA SP-483.
- Galdes, D. and Smith, P., Ohio State U., personal communication, March, 1990.
- Harrison, Albert A. and Connors, Mary M. (1984) Groups in Exotic Environments. Advances in Experimental Social Psychology, 18, L. Berkowitz (ed), Academic Press, pp. 50-87.
- Norman, Donald A. (1987) Cognitive Engineering-Cognitive Science. Chapter 12 in J.M. Carroll (ed) Interfacing Thought. Cambridge, Mass: A Bradford Book, The MIT Press.
- Norman, Donald A. (1990) The Problem with Automation: Inappropriate Feedback and Interaction, Not "Over-automation". Phil. Trans. R. Soc. Lond. B 000, 000-000.
- Statler, Irving C. and Billings, Charles E. (1989) Maintaining Human Productivity During Mars Transit. Presented at the 19th Intersociety Conference on Environmental Systems, 24-26 July 1989, San Diego. SAE Technical Paper Series 891435.
- Wiener, Earl L. and Nagel, David C. (eds.) (1988) Human Factors in Aviation San Diego: Academic Press, Inc.
- Winograd, Terry and Fernando, Flores (1987) Understanding Computers and Cognition. Addison-Wesley Publishing Co.